

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 788 280 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
15.04.1998 Bulletin 1998/16

(51) Int Cl.⁶: **H04N 7/173**

(43) Date of publication A2:
06.08.1997 Bulletin 1997/32

(21) Application number: **97300405.4**

(22) Date of filing: **22.01.1997**

(84) Designated Contracting States:
DE FR GB

(30) Priority: **01.02.1996 US 595313**

(71) Applicant: **INTERNATIONAL BUSINESS
MACHINES CORPORATION**
Armonk, NY 10504 (US)

(72) Inventors:
• **Aggarwal, Charu Chandra,**
Mass. Inst. of Technology
Cambridge, Massachusetts 02139 (US)

• **Wolf, Joel Leonard**
Katonah, New York 10536 (US)
• **Yu, Philip Shi-Lung**
Chappaqua, New York 10514 (US)

(74) Representative: **Moss, Robert Douglas**
IBM United Kingdom Limited
Intellectual Property Department
Hursley Park
Winchester Hampshire SO21 2JN (GB)

(54) **Maximum factor selection policy for batching VOD requests**

(57) A VOD scheduler maintains a queue of pending performance for each video. Using the notion of queue selection factor, a batching policy is devised that schedules the video with the highest selection factor. Selection factors are obtained by applying discriminatory weighting factors to the adjusted queue lengths associated with each video where the weight decreases as the popularity of the respective video increases and the queue length is adjusted to take defection into account.

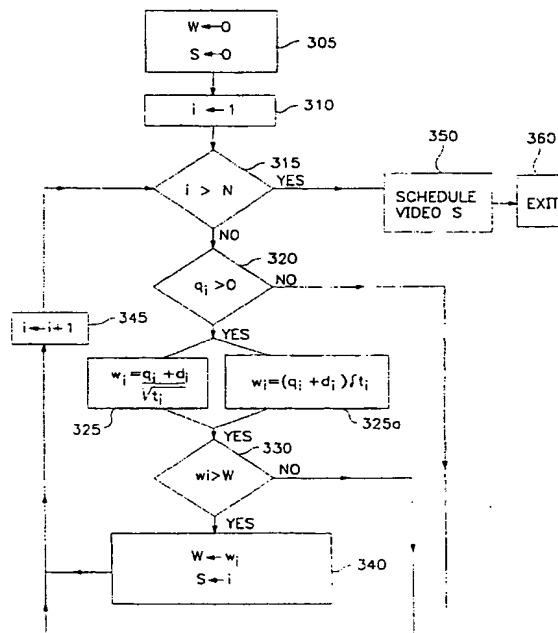


FIG. 3

EP 0 788 280 A3

This Page Blank (uspto)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 97 30 0405

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	US 5 351 075 A (HERZ FREDERICK ET AL) * column 2, line 50 - column 4, line 25 *	1,2,6,7	H04N7/173 G11B27/00
A	PATENT ABSTRACTS OF JAPAN vol. 007, no. 212 (P-224), 20 September 1983 & JP 58 107962 A (FUJITSU KK), 27 June 1983, * abstract *	1,6	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			H04N G11B
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 25 February 1998	Examiner Fuchs, P
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03 82 (P04C01)

This Page Blank (uspto)

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 788 280 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
06.08.1997 Bulletin 1997/32

(51) Int Cl.⁶: H04N 7/173

(21) Application number: 97300405.4

(22) Date of filing: 22.01.1997

(84) Designated Contracting States:
DE FR GB

(30) Priority: 01.02.1996 US 595313

(71) Applicant: INTERNATIONAL BUSINESS
MACHINES CORPORATION
Armonk, NY 10504 (US)

(72) Inventors:
• Aggarwal, Charu Chandra,
Mass. Inst. of Technology
Cambridge, Massachusetts 02139 (US)

• Wolf, Joel Leonard
Katonah, New York 10536 (US)
• Yu, Philip Shi-Lung
Chappaqua, New York 10514 (US)

(74) Representative: Moss, Robert Douglas
IBM United Kingdom Limited
Intellectual Property Department
Hursley Park
Winchester Hampshire SO21 2JN (GB)

(54) Maximum factor selection policy for batching VOD requests

(57) A VOD scheduler maintains a queue of pending performance for each video. Using the notion of queue selection factor, a batching policy is devised that schedules the video with the highest selection factor. Selection

factors are obtained by applying discriminatory weighting factors to the adjusted queue lengths associated with each video where the weight decreases as the popularity of the respective video increases and the queue length is adjusted to take defection into account.

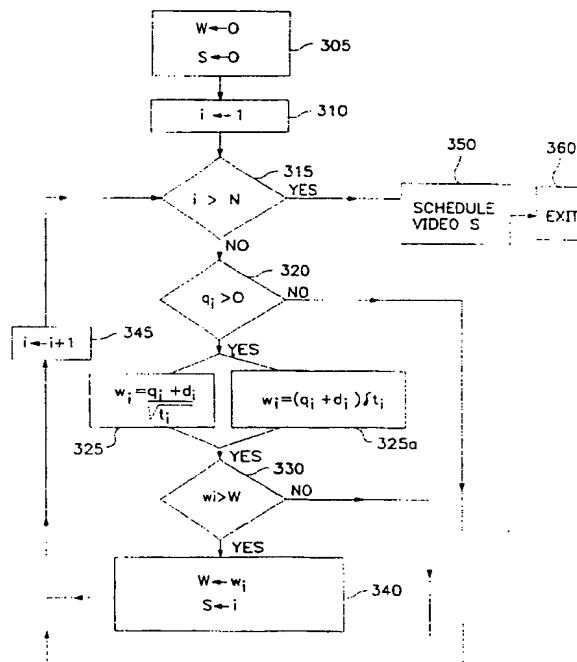


FIG. 3

Description

BACKGROUND OF THE INVENTION

The present invention relates to the scheduling of performance requests in Video-On-Demand (VOD) systems.

In Video-on-Demand systems, popular videos (i.e. movies or other programs) are often requested by many viewers. In order to increase system throughput, some viewers requesting an identical video may be transmitted (and hence share) a single video stream. This is referred to in the art as batching.

Depending upon the system load, a requested video may not be started immediately. While viewers will usually tolerate a small wait time, a long wait may result in the loss (or defection) of viewers. This may be accomplished, for example, by the viewer indicating (either actively or passively) that he is no longer interested in viewing the requested video.

A conventional approach to video scheduling is the use of a first-come-first-serve (FCFS) policy. Under this policy, all video requests are placed in a single request queue. The request at the front of the queue is served when system capacity is available to service the request. If batching is supported, all subsequent queued requests for the same video as the serviced request are also serviced from the same signal stream.

An alternative to the above approach is to maintain a separate request queue for each video and to select the video with the longest queue for the next showing. This is referred to as the longest queue length first (LQF) policy. Still another approach is to provide periodic showing (say every 5 minutes) for the most popular videos. For requests which would not be serviced by periodic showings, another scheduling scheme such as FCFS can be used.

As mentioned above, a viewer with a waiting request may defect if the waiting time exceeds the viewer's tolerance. Viewer defection is, of course, undesirable. The choice of video scheduling policy can have a significant effect on the amount of batching and defection. The FCFS policy does not take into account the batch size. By contrast, the LQF policy ignores the wait time already incurred by the waiting requests. Periodic showing (when used alone) may not provide the flexibility needed to cope with dynamic load variations.

Recently, a new approach on batching was proposed in Shachnai, H., and Yu, P., The Role of Wait Tolerance in Effective Batching: A Paradigm for Multimedia Scheduling Schemes, IBM Research Report RC 20038, April 1995, which is incorporated herein by reference. Under this approach, the additional wait time that viewers will tolerate before defecting, in addition to a video's respective queue length, is considered in making scheduling decisions. When system capacities become available, rather than scheduling the video immediately, the scheduler delays transmission of the video until just

prior to expiration of the maximum wait tolerance time of the longest waiting one of the pending transmission requests. In the interim, additional requests for transmission of the video may join the queue. In contrast with LQF, a video with the longest queue may not get scheduled if all waiting requests are fairly recent. This prevents the most popular videos from monopolizing the stream capacities. In contrast with FCFS, better batching is achieved. To effectively implement this algorithm, however, knowledge of viewer tolerance is required. This knowledge may not be readily available.

SUMMARY OF THE INVENTION

According to a first aspect, the present invention provides a method of scheduling a plurality of video requests, comprising the steps of: placing each of said plurality of video requests into a respective corresponding queue; determining a queue selection factor for each queue based on length of each queue and defections in each queue; and scheduling one of said plurality of video requests based on said queue selection factor of each of said queues.

According to a second aspect, the present invention provides video apparatus for scheduling a plurality of video requests, comprising: means for receiving a plurality of video requests and for maintaining a plurality of queues, each of said queues storing common ones of said requests means for determining a queue selection factor for each queue based on length of each queue and defections in each queue; and means for scheduling one of said plurality of video requests based on said queue selection factor of each of said queues.

A VOD scheduler maintains a queue of pending performance for each video. Using the notion of queue selection factor, a batching policy is devised that schedules the video with the highest selection factor. Selection factors are obtained by applying discriminatory weighting factors to the adjusted queue lengths associated with each video where the weight decreases as the popularity of the respective video increases and the queue length is adjusted to take defection into account.

Two alternative means of obtaining queue selection factors are provided. The first approach, referred to as the steady state approach, is designed for quasi static or steady state environments where the relative request frequencies of the different videos changes very gradually (in terms of hours). The second approach, referred to as the instantaneous approach, is for a more dynamic environment where the relative request frequencies of the different videos change rapidly (in terms of minutes).

Under the steady state approach, the queue selection factor of a video is obtained by dividing the adjusted queue length of the video's associated queue by the square root of its relative request frequency. To account for the effect of defection, the queue length is adjusted by the number of defections. That is to say that the scheduler substitutes the queue length in the selection

factor calculation by the number of requests for a video since the last time that it was scheduled. Thus, for example, if a large percentage of the viewers corresponding to unpopular videos defect anyway, then using the number of requests since last scheduling (instead of the queue length) increases the fairness of the system.

One way of estimating the relative frequency is to use a moving window estimate. This is obtained by counting, for each video, the number of requests which had occurred in a moving window of the immediate past (for example, a window of the past hour). Smoothing techniques can also be applied to take weighted averages over past estimates.

Since unpopular videos have respectively lower relative request frequencies, the square roots of the inverse of their relative frequencies are larger than those of the popular videos. Thus, the queue selection factors are obtained by applying discriminatory weighting factors to the adjusted queue lengths which serves as a bias against the more popular videos. This prevents popular videos from monopolizing all the stream capacities and forming frequent batches of small sizes (which is what occurs using the LQF policy).

Under the instantaneous approach, the queue selection factor of a video is obtained by multiplying its adjusted queue length by the time which has elapsed since the last scheduling of the video. Although unpopular videos have lower request frequencies or shorter queues, their selection factors will increase as the time elapsed since last scheduling increases. This again provides discriminatory weighting factors which disfavour the more popular videos.

DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a multimedia server;

FIGs. 2A-2B are flow charts of events handled in the Maximum Selection Factor (MSF) scheduler; and

FIG. 3 is a flow chart of a video request handled by the MSF scheduler.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Figure 1 is a block diagram of a video server 100 in which video data (movies or other moving visual or audio visual performances) is stored in disks 102 and is transmitted to the end client stations 103 over network 104 upon request. The video server 100 also includes disks (not shown) which store working data and the program code for the video server 100. Other storage devices such as video tape players and jukeboxes (from which movies can be loaded to disks 102) may be optionally included.

The program code for the video server 100 includes

processes such as the server main control program, a video scheduling program, a customer usage tracking program, and conventional communications, I/O and buffer management processes. The video server 100 includes a processor (i.e. CPU) 101 which executes the processes under control of the main control program and a memory buffer 105 which is used for the temporary storage of portions of videos (thus enabling some users to be served from the memory buffer 105 rather than the disks 102). The memory buffer 105 can also be used to handle pause/resume requests by providing a temporary storage area for portions of a video which stream from the disk while a viewer is paused. Other than the scheduler 106, the process may be of a conventional type typically used in VOD systems and which is described, for example in Yu, P. et al., Design and analysis of a look-ahead scheduling scheme to support pause-resume for video-on-demand applications, Multimedia Systems, (1995); and Wolf, J. et al., DASH: A Disk Load Balancing Optimization Scheme for Video-On-Demand Computer Systems ACM SIGmetrics, Ottawa, Canada (1994).

The video server 100 may be embodied using any processor of sufficient performance for the number of video streams to be supported. For example, a small capacity video server may be embodied using a RISC System/6000 TM system while a larger capacity server could be embodied using an ES/9000 TM system (both available from International Business Machines Corporation of Armonk, New York). This disk 102 may be any conventional disk or disk array. The communication network 104 may be, for example, a fibre optic network or a conventional bidirectional cable network. The client stations 103 may each be embodied as a set-top box.

The scheduler 106 performs a number of tasks which result in the scheduling of performances for requested videos. In accordance with an exemplary embodiment of the present invention, upon receiving video requests from the client station 103, the scheduler (maximum selection factor based) 106 enters each request into the appropriate video queue. It tracks the last scheduling time of each video, the number of defections since last scheduling, and the relative request frequency for each video.

Assume that there are N different videos and a separate queue is maintained for each video. Let q_i be the queue length of the i -th video queue, d_i be the number of defections since last scheduling at the i -th video queue, f_i be the relative request frequency of the i -th video, and δt_i represents the time since last scheduling of video i .

The present invention employs the notion of queue selection factor and devises the MSF policy that schedules the video with the maximum selection factor. The queue selection factor is obtained by applying discriminatory weighting factors to the respective queue lengths of the different videos where the weighting decreases as the popularity of the video increases.

Two alternative means of obtaining queue selection factor are provided. The steady state approach is designed for a quasi static or steady state environment where the relative request frequencies of the different videos change very gradually (in terms of hours). The instantaneous approach is for a more dynamic environment where the relative request frequencies of the different videos can change rapidly (in terms of minutes).

Under the steady state approach, the queue selection factor of a video is obtained by dividing the adjusted queue length of a video s queue by the square root of the video s relative request frequency. The adjusted queue length of a video s queue is the number of requests since last scheduling, the queue length adjusted by the number of defections since last scheduling. w_i , the queue selection factor of the i -th video queue is obtained as follows:

$$w_i = \frac{q_i + d_i}{\sqrt{f_i}}$$

One method of estimating the relative request frequency is to use the moving window estimate which is obtained by counting the number of requests for each video which occur in a moving window of a past period (for example a window of the past hour). Smoothing techniques for forecasting can also be applied to take weighted average over past estimates (as described for example by F. Hiller and O. Lieberman in Operations Research, Second Edition (1974), Holden-day, pages 522-526).

Since unpopular videos have lower relative request frequencies, the square roots of the inverse of their relative frequencies are larger than those of the popular videos. Thus, the queue selection factors are obtained by applying discriminatory weighting factors to the adjusted queue lengths which disfavour against the more popular videos. This would prevent the popular videos from monopolizing all the stream capacities and forming frequent batches of small sizes (which is what occurs using the LQF policy).

Under the instantaneous approach, the queue selection factor of a video is obtained by multiplying its adjusted queue length of the video s queue by the time elapsed since last scheduling of the video. This is expressed as:

$$w_i = (q_i + d_i) \delta t_i$$

Although unpopular videos have lower request frequencies or shorter queues, their selection factors increase as the time elapsed since last scheduling increases. This again provides discriminatory weighting factors which disfavour the more popular videos.

The implication of the MSF scheduling policy on the fairness of the system compared to the LQF policy is next examined. Intuitively, the fairness of the system measures the variation in the quality of service provided across different videos. LQF is rather unfair to the unpopular videos, because the queue for a popular video tends to build up much faster than the queue for an unpopular video. Using the queue selection factor evens out the unfairness by favouring the unpopular videos. Thus, MSF is much fairer than the LQF policy.

FIGS. 2A-2B are flow charts of events handled in the MSF schedule. Each of the tasks can execute in parallel for different requests. However, the video stream scheduler is desirably a serialization point. That is to say the video stream scheduler can be preferably invoked by only one task or event at a time and executes to completion once invoked.

FIG. 2A is a flow chart of video request handling by the MSF scheduler. Each time a new performance request for a video is made by a viewer, its arrival at the VOD system is detected by the scheduler at step 202. Next, at step 204, the scheduler determines the appropriate queue for the request based on the video requested and enters the request into the correct queue accordingly. Then, at step 206, the scheduler determines if there is any stream capacity available on the server. If there is no capacity available to service the request, the scheduler exits at step 208. At this point the request can not be scheduled until a currently running video completes and its associated stream capacity is freed. If the server has a stream capacity available, then at step 210 the scheduler invokes the video stream scheduling task of FIG. 3.

FIG. 2B is a flow chart of video completion handled by the MSF scheduler. Completion of a video is detected by the scheduler at step 212. A video can complete either by finishing the performance through its scheduled end time or by all of the viewers of the video exiting the video (terminating the performance). At step 214, the scheduler frees the stream capacity by marking it as being available" (the use or available status of each stream capacity is conventionally tracked by the video server by way of a status table).

At step 216 the scheduler checks the request queues to determine if they are empty (i.e. there are no pending requests on any of the video queues). If the request queues are all empty, at step 218 the scheduler exits. Otherwise, if the request queues are not empty (there are waiting requests), at step 220 the scheduler invokes the video stream scheduling task of FIG. 3.

FIG. 3 is a flow chart of video requests handled by the MSF scheduler to calculate queue selection factor. The scheduler 106 is invoked when a stream becomes available or a new video request arrives. At step 305, W and S are set equal to zero, where W tracks the maximum selection factor among the video queues scanned and S tracks the corresponding index of the video queue. At step 310, i is initialized to 1, where i is the

running index to scan through all video queues.

At step 315, the scheduler checks if the value of i is larger than the number N of video queues. This step is performed to determine whether all video queues have been examined. If not, the scheduler proceeds to step 320. At step 320, if q_i is not zero, the queue selection factor (w_i) is calculated either at step 325 or 325a (depending on whether the steady state approach or the instantaneous approach is used). At step 330, w_i is compared to W . If w_i is larger, at step 340 the scheduler resets the value of W to track the largest queue selection factor scanned so far and also S to indicate that the i -th video queue has the largest queue selection factor among the first i video queues. At step 345, i is incremented by 1. Step 315 is then performed.

In accordance with the description of step 315 set forth above, if the value of i is larger than N , then all video queues have been examined. At step 350, the scheduler selects video S to schedule and then exits at step 360.

While the invention has been described in terms of exemplary embodiments, it is contemplated that it may be practical as outlined above with modifications within the spirit and scope of the appended claims.

Claims

1. A method of scheduling a plurality of video requests, comprising the steps of:

a) placing each of said plurality of video requests into a respective corresponding queue;

b) determining a queue selection factor for each queue based on length of each queue and defections in each queue; and

c) scheduling one of said plurality of video requests based on said queue selection factor of each of said queues.

2. A method of scheduling a plurality of video requests according to claim 1, wherein said queue selection factor is based on video request frequency of each of said plurality of video requests.

3. A method of scheduling a plurality of video requests according to claim 1, wherein said queue selection factor is based on time since last scheduling in said respective corresponding queue.

4. A method of scheduling a plurality of video requests according to claim 2, wherein for queue i :

q_i = queue length
 d_i = defections since last scheduling
 f_i = relative request frequency

and the queue selection factor is:

$$\frac{q_i + d_i}{\sqrt{f_i}}$$

5. A method of scheduling a plurality of video requests according to claim 3, wherein for queue i :

q_i = queue length
 d_i = defections since last scheduling
 δt_i = time since last scheduling

and the queue selection factor is:

$$(q_i + d_i) \delta t_i$$

6. Video apparatus for scheduling a plurality of video requests, comprising:

means for receiving a plurality of video requests and for maintaining a plurality of queues, each of said queues storing common ones of said requests

means for determining a queue selection factor for each queue based on length of each queue and defections in each queue; and

means for scheduling one of said plurality of video requests based on said queue selection factor of each of said queues.

7. Apparatus for scheduling a plurality of video requests according to claim 6 wherein said queue selection factor is based on video request frequency of each of said plurality of video requests.

8. Apparatus for scheduling a plurality of video requests according to claim 6 wherein said queue selection factor is based on time since last scheduling in said respective corresponding queue.

9. Apparatus for scheduling a plurality of video requests according to claim 7, wherein for queue i :

q_i = queue length
 d_i = defections since last scheduling
 f_i = relative request frequency

and the queue selection factor is:

$$\frac{q_i + d_i}{\sqrt{f_i}}$$

10. Apparatus for scheduling a plurality of video requests according to claim 8, wherein for queue i :

q_i = queue length

d_i = defections since last scheduling

δt_i = time since last scheduling

and the queue selection factor is:

5

$(q_i + d_i) \delta t_i$.

10

15

20

25

30

35

40

45

50

55

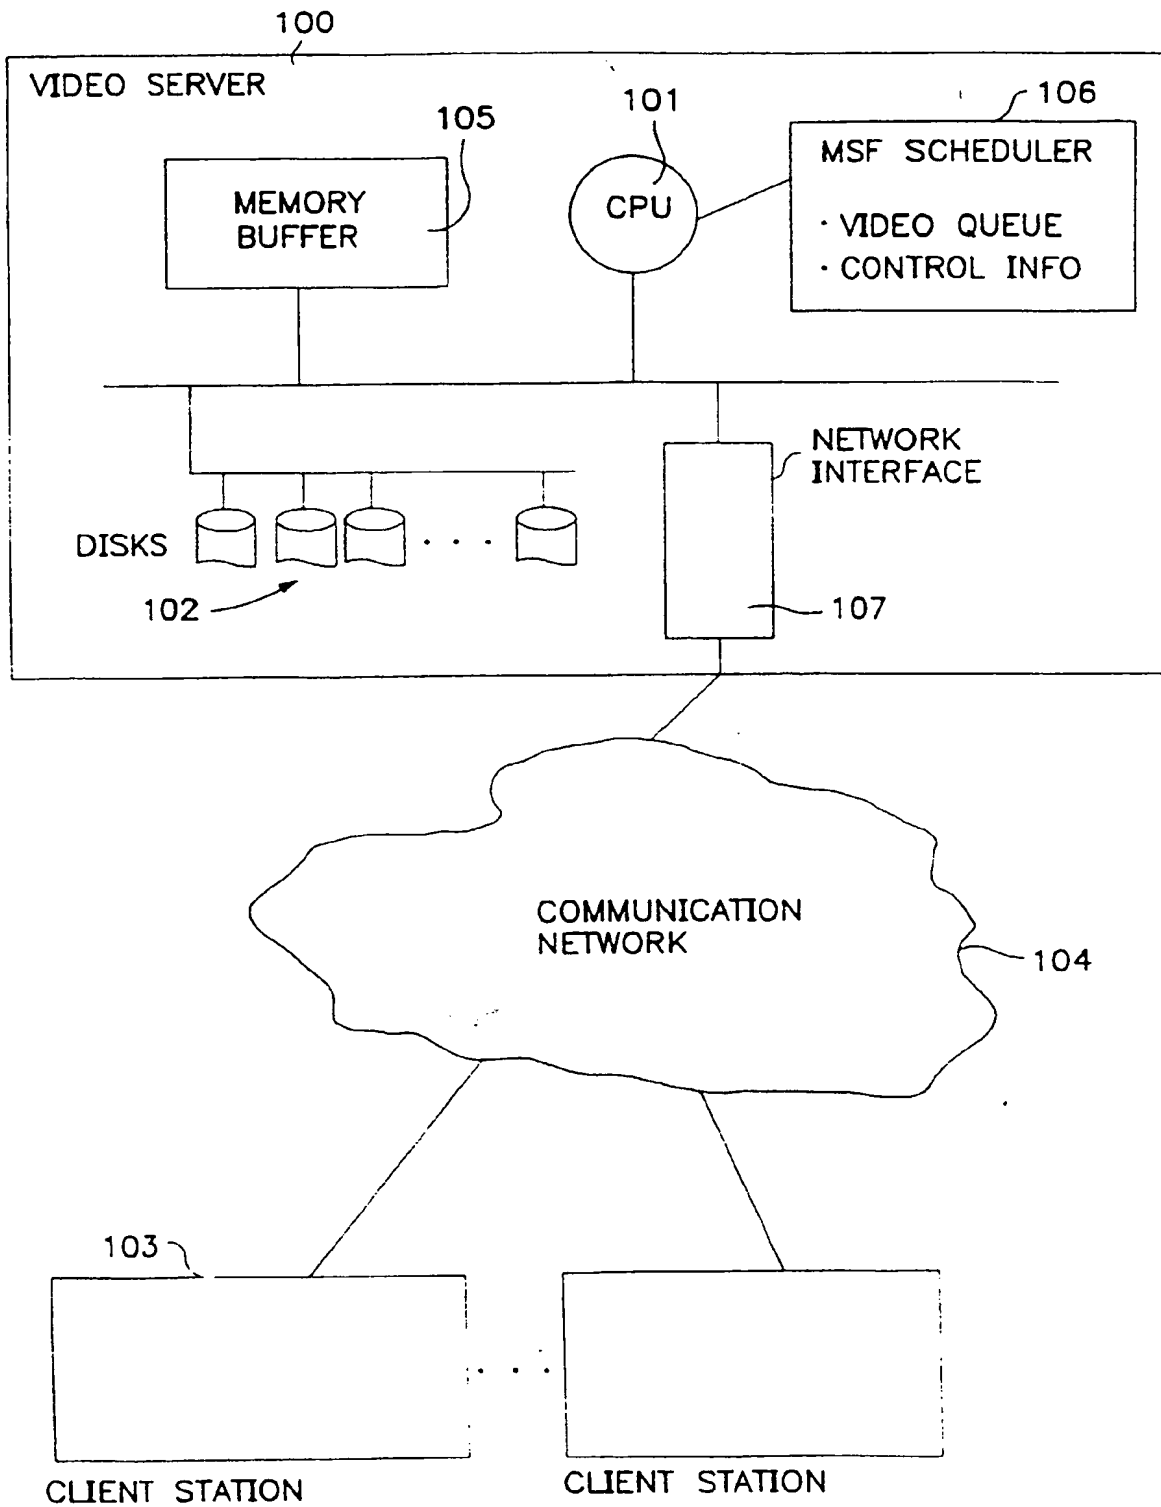


FIG. 1

This Page Blank (uspto)

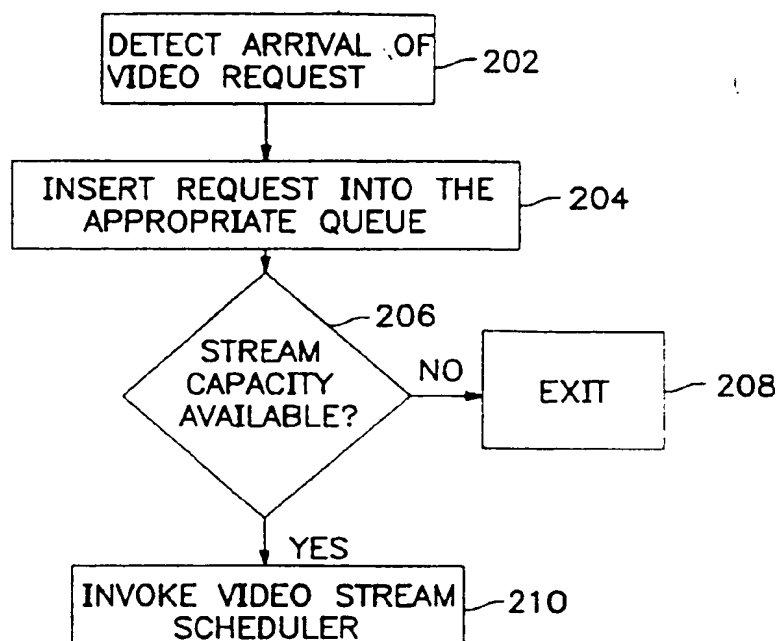


FIG. 2A

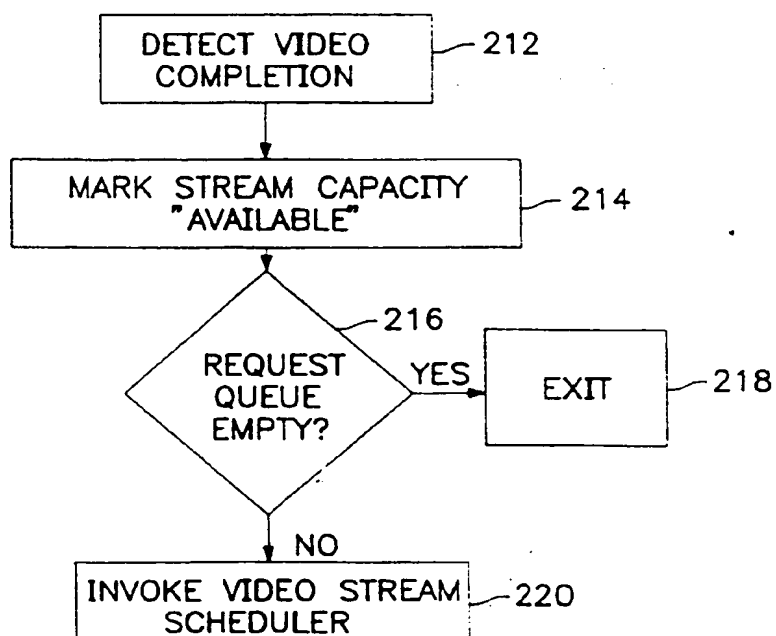


FIG. 2B

This Page Blank (uspto)

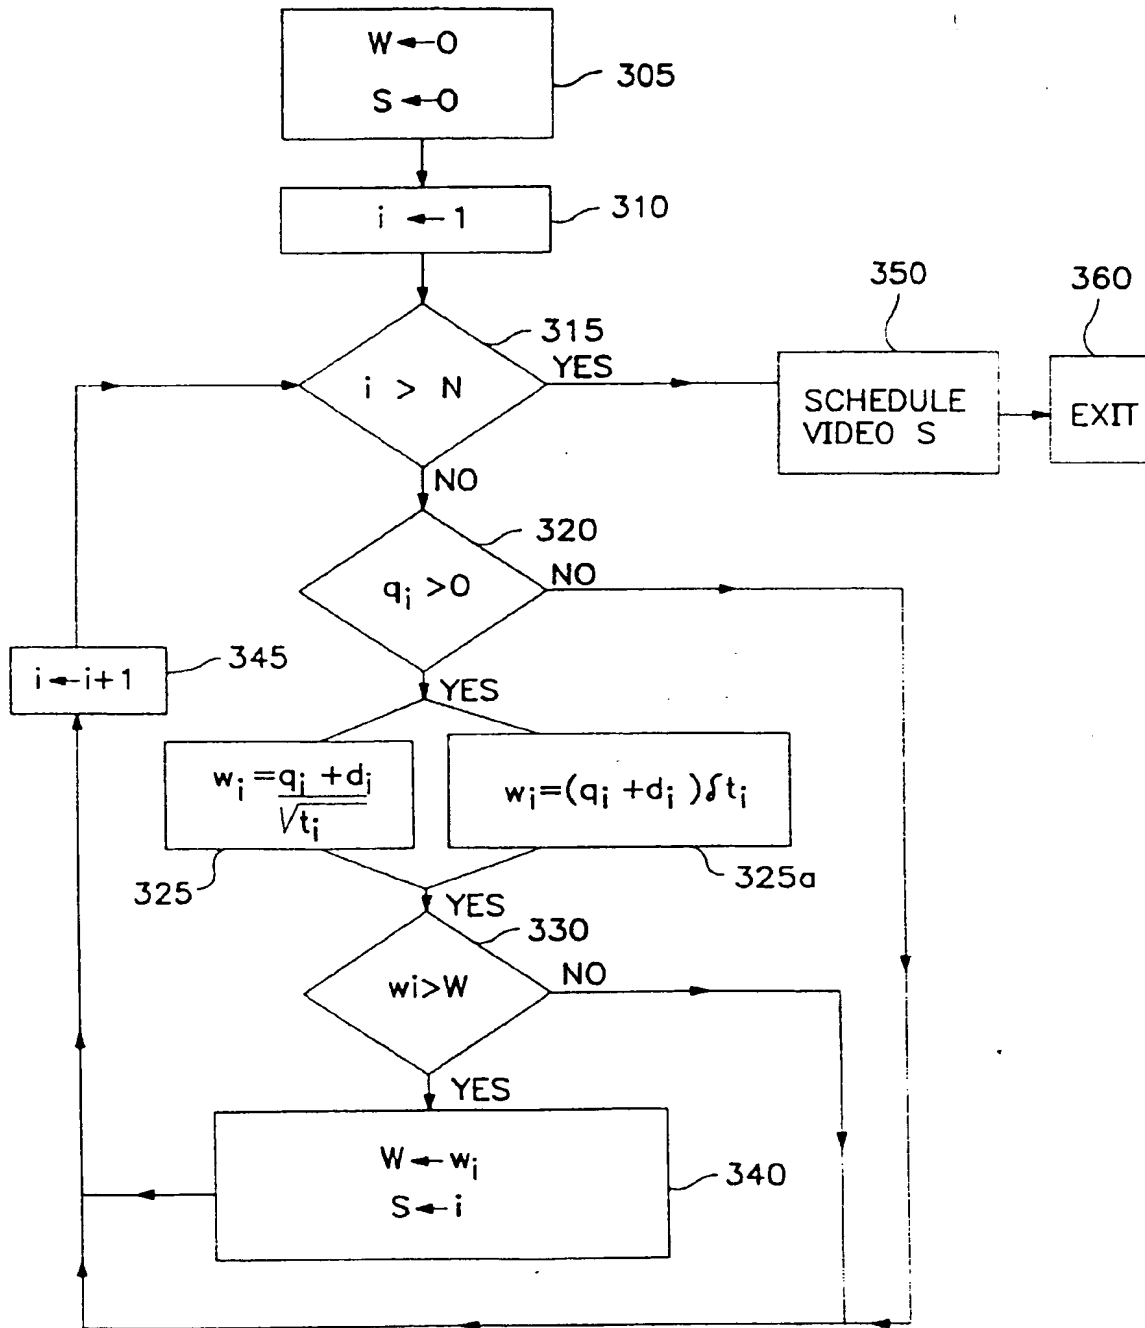


FIG. 3

This Page Blank (uspto)